

CLAIMS

- 1 1. A method for taking over a failed filer owning disks that store file service data
2 and owning at least one disk that is free of file service data by a clustered partner filer,
3 the failed filer being adapted to perform a coredump in which, in the event of failure,
4 memory contents of the failed filer are transferred to a disk, the method comprising the
5 steps of:
6 changing, by the failed filer, a coredump attribute on the disk that is free of file
7 service data (the “coredump disk”) from a non-coredump state to a coredump state and
8 maintaining the coredump attribute on other disks owned by the failed filer in a non-
9 coredump state;
10 writing the memory contents to the coredump disk;
11 identifying, by the clustered partner filer, the coredump attribute of the other disks
12 and taking ownership of the other disks while allowing the failed filer to maintain owner-
13 ship of the coredump disk;
14 upon completion of the writing of the memory contents, changing the coredump
15 attribute to a non-coredump state; and
16 upon identification of the non-coredump state in the coredump attribute of the
17 coredump disk, taking ownership, by the clustered partner filer, of the coredump disk.
- 1 2. The method as set forth in claim 1 where the coredump disk is a spare disk owned
2 by the failed filer.
- 1 3. The method as set forth in claim 1 where the coredump disk is a disk dedicated to
2 a coredump.
- 1 4. The method as set forth in claim 1 further comprising the step of creating, from
2 the memory contents written on the coredump disk, a coredump data set for diagnosis of
3 the failing filer and writing the coredump file to the file system root of the failing filer.

1 5. The method as set forth in claim 1 where the step of writing the memory contents
2 includes limiting writing to the coredump disk to a predetermined time limit following a
3 failure event of the failed filer.

1 6. The method as set forth in claim 5 wherein the predetermined time limit is less
2 than a maximum time a panic of a network communicating with each of the failed filer
3 and the clustered partner filer occurs.

1 7. The method as set forth in claim 5 wherein the step of limiting includes, upon
2 elapsing of the predetermined time limit, changing the coredump attribute of the
3 coredump disk to a non-coredump state.

1 8. The method as set forth in claim 1 wherein the coredump disk and each of the
2 other disks owned by the failed filer include a plurality of predetermined regions includ-
3 ing a coredump region and a file system data storage region and wherein the step of
4 writing of the memory contents includes writing the memory contents into the file system
5 data storage region of the coredump disk.

1 9. The method as set forth in claim 8 wherein the coredump region includes a
2 coredump header adapted to store the coredump attribute.

1 10. A storage system including a first server owning interconnected first storage de-
2 vices and a second server owning interconnected second storage devices, the first server
3 and the second server being connected together by a cluster interconnect so that the sec-
4 ond server can take over ownership of the first storage devices upon failure of the first
5 server; the storage system comprising:

6 a coredump function that (a) causes the first server to write its memory to a
7 coredump storage device chosen from one of the first storage devices in response to a
8 sensed failure of the first server, each of the first storage devices including a coredump
9 attribute (b) causes the attribute of the coredump storage device to be set to a coredump

10 state and the coredump attribute of other of the first storage devices to be set to a non-
11 coredump state; and

12 a takeover function that (a) identifies each of the first storage devices with the
13 coredump attribute set to the non-coredump state, (b) changes of each of the second de-
14 vices having the coredump attribute set to the non-coredump state from ownership by the
15 first server to ownership by the second server so that takeover of the ownership can pro-
16 ceed in parallel with the writing of the memory to the coredump storage device.

1 11. The storage system as set forth in claim 10 wherein the takeover function is
2 adapted to cause the second server to logically assume ownership of the coredump stor-
3 age device after the earlier of either an elapsing of a predetermined time limit or a com-
4 pletion of the writing of the memory to the coredump storage device.

1 12. The storage system as set forth in claim 11 wherein the coredump function is
2 adapted to cause the coredump attribute of the coredump storage device to be changed to
3 a non-coredump state after either the elapsing of the predetermined time limit or the
4 completion of the writing of the memory to the coredump storage device.

1 13. The storage system as set forth in claim 12 wherein the non-coredump state com-
2 prises each of an aborted state, a completed state and a non-active state.

1 14. The storage system as set forth in claim 10 wherein the first storage devices each
2 comprise a disk drive engaged in file service activity and the coredump storage device
3 comprises a spare disk that is free of file service activity.

1 15. A computer-readable medium in a storage system that includes a first server
2 owning interconnected first storage devices and a second server owning interconnected
3 second storage devices, the first server and the second server being connected together by
4 a cluster interconnect so that the second server can take over ownership of the first stor-

5 age devices upon failure of the first server; the computer-readable medium including pro-
6 gram instructions for performing the steps of:

7 writing, by the first server, its memory to a coredump storage device chosen from
8 one of the first storage devices in response to a sensed failure of the first server, each of
9 the first storage devices including a coredump attribute;

10 setting the attribute of the coredump storage device to be set to a coredump state
11 and the coredump attribute of other of the first storage devices to be set to set to a non-
12 coredump state;

13 identifying each of the first storage devices with the coredump attribute set to the
14 non-coredump state;

15 changing of each of the second devices having the coredump attribute set to the
16 non-coredump state from ownership by the first server to ownership by the second server
17 so that takeover of the ownership can proceed in parallel with the writing of the memory
18 to the coredump storage device.

1 16. The computer-readable medium as set forth in claim 15 further comprising the
2 step of causing the second server to logically assume ownership of the coredump storage
3 device after the earlier of either an elapsing of a predetermined time limit or a completion
4 of the writing of the memory to the coredump storage device.

1 17. The computer-readable medium as set forth in claim 16 further comprising the
2 step of causing the coredump attribute of the coredump storage device to be changed to a
3 non-coredump state after either the elapsing of the predetermined time limit or the com-
4 pletion of the writing of the memory to the coredump storage device.

1 18. The computer-readable medium as set forth in claim 17 wherein the non-
2 coredump state comprises each of an aborted state, a completed state and a non-active
3 state.

1 19. The computer-readable medium as set forth in claim 18 wherein the first storage
2 devices each comprise a disk drive engaged in file service activity and the coredump
3 storage device comprises a spare disk that is free of file service activity.

1 20. The computer-readable medium as set forth in claim 15 wherein each of the first
2 storage devices includes a coredump information region and a file system region and
3 wherein the memory is written into the file system region of the coredump storage de-
4 vice.

1 21. The computer-readable medium as set forth in claim 15 further comprising the
2 step of creating, with the second server, a coredump data set from the memory written to
3 the coredump storage device, the data set being adapted to enable diagnosis of a fault
4 relative to the first server.

1 22. The computer-readable medium as set forth in claim 21 further comprising the
2 step of writing the coredump data set by the second server to a file system root of the first
3 server stored on the first storage devices.

1 23. A method for takeover in a storage system that includes a first server owning in-
2 terconnected first storage devices and a second server owning interconnected second
3 storage devices, the first server and the second server being connected together by a
4 communication interconnect so that the second server can take over ownership of the first
5 storage devices upon failure of the first server, the method comprising the steps of:

6 writing, by the first server, its memory to a coredump storage device chosen from
7 one of the first storage devices in response to a sensed failure of the first server, each of
8 the first storage devices including a coredump attribute;

9 setting the attribute of the coredump storage device to be set to a coredump state
10 and the coredump attribute of other of the first storage devices to be set to set to a non-
11 coredump state;

12 identifying each of the first storage devices with the coredump attribute set to the
13 non-coredump state; and

14 changing each of the second storage devices having the coredump attribute set to
15 the non-coredump state from ownership by the first server to ownership by the second
16 server so that takeover of the ownership can proceed in parallel with the writing of the
17 memory to the coredump storage device.

1 24. The method as set forth in claim 23 further comprising the step of causing the
2 second server to logically assume ownership of the coredump storage device after the
3 earlier of either an elapsing of a predetermined time limit or a completion of the writing
4 of the memory to the coredump storage device.

1 25. The method as et forth in claim 24 wherein the step of changing includes setting a
2 reservation on each of second storage devices so as to establish ownership by the second
3 server.

1 26. The method as set forth in claim 25 wherein the reservation comprises a SCSI-3
2 reservation.